The Ensemble Kalman Filter for Inverse Problems

Marco A. Iglesias∗, Kody J.H. Law∗ and Andrew M. Stuart∗

∗ Mathematics Institute, University of Warwick, Coventry CV4 7AL, UK

E-mail: {M.A.Iglesias-Hernandez,K.J.H.Law,A.M.Stuart}@warwick.ac.uk

Abstract.

The Ensemble Kalman filter (EnKF) was introduced by Evensen in 1994 [1] as a novel method for *data assimilation*: state estimation for noisly observed time-dependent problems. Since that time it has had enormous impact in many application domains because of its robustness and ease of implementation, and numerical evidence of its accuracy. In this paper we demonstrate that the EnKF can be generalized to a wide class of *inverse problems*. In this context we show that the the EnKF estimate of the unknown function lies in a subspace $\mathcal{A}$ spanned by the initial ensemble. Hence the resulting error may be bounded above by the error found from best approximation in this subspace. We provide numerical experiments which compare the error incurred by EnKF for inverse problems with the error of the best approximation in $\mathcal{A}$, and with variants on traditional least-squares approaches, restricted to the subspace $\mathcal{A}$. In so doing we demonstrate that the EnKF method for inverse problems provides a derivative-free optimization method with comparable accuracy to that achieved by traditional least-squares approaches. Furthermore, we also demonstrate that accuracy is of the same order of magnitude as that achieved by best approximation. Three examples are used to demonstrate these assertions: inversion of a compact linear operator; inversion of piezometrix head to determine hydraulic conductivity in a Darcy model of groundwater flow; and inversion of Eulerian velocity measurements at positive times to determine the initial condition in an imcompressible fluid.

Submitted to: *Inverse Problems*

1. Introduction

Since its introduction in [1], the Ensemble Kalman filter (EnKF) has had enormous impact on applications of data assimilation to state estimation, and in particular in oceanography [2], reservoir modelling [3] and weather forecasting [4]; the books [5, 6, 7] give further details and references to applications in these fields. However, although demonstrably accurate, robust and simple to implement, the EnKF is not widely analyzed and its properties are poorly understood. The purpose of this paper is threefold: (i) to demonstrate that a novel non-standard perspective on the EnKF methodology converts it into a generic tool for solving inverse problems; (ii) to provide some basic analysis of the properties of this algorithm for inversion; (iii) to demonstrate numerically that the method can be effective on a wide range of applications.

In section 2 we describe the general inverse problem framework within which we work. We highlight the fact that the EnKF approach to this inverse problem provides an approximation to the solution in the subspace $\mathcal{A}$ spanned by the initial ensemble members, a fact observed for a specific sequential implementation of the EnKF in [8]. We then describe two classes of algorithms which we will use to evaluate the EnKF methodology. The first is best approximation of the truth in the subspace $\mathcal{A}$. This is, of course, not an implementable method as the truth is not known; however it provides an important lower bound on the achievable error of the EnKF method and hence has a key conceptual role. The second class are variants on least squares, restricted to the subspace $\mathcal{A}$. In section 3 we describe the EnKF method for the general inverse problem, which is in the form of a nonlinear iteration of an interacting ensemble. In Theorem 3.1 we prove that the method gives rise to an approximation which lies in the subspace $\mathcal{A}$ spanned by the initial ensemble members and Corollary 3.3 uses this fact to give a lower bound on the achievable approximation error of the EnKF algorithm.

Section 4 contains numerical experiments which illustrate the ideas in this paper on a linear inverse problem. The forward operator is a compact operator found from inverting the negative Laplacian plus identity. In section 5 we numerically study the groundwater flow inverse problem of determining hydraulic conductivity from piezometric head measurement in an elliptic Darcy flow model. Section 6 contains numerical results concerning the problem of determining the initial condition for the velocity in a Navier-Stokes model of an incompressible fluid; the observed data are pointwise (Eulerian) measurements of the velocity field.

The numerical results in this paper all demonstrate that the EnKF method for inversion is a derivative-free regularized optimization technique which produces numerical results similar in accuracy to those found from least-squares based methods in the same subspace $\mathcal{A}$. Furthermore, the three examples serve to illustrate the point that the method offers considerable flexibility through the choice of initial ensemble, and hence the subspace $\mathcal{A}$ in which it produces an approximation. In particular, for the linear and Darcy inverse problems, we make two choices of initial ensemble: (i) draws from a prior Gaussian measure and (ii) the Karhunen-Loéve basis functions of the centered Gaussian measure shifted by the mean. For the Navier-Stokes inverse problem the initial ensemble is also chosen to comprise randomly drawn functions on the attractor of the dynamical system.

2. The Problem

2.1. The Set-Up

Let $\mathcal{G} : X \to Y$ where X and Y are Hilbert spaces. In the following, we use $\langle \cdot, \cdot \rangle$ and $\| \cdot \|$, etc. as the inner-product and norm on both X and Y, and it will be clear from the context which space is intended. Also, we let $\| \cdot \|_A$ denote $\|A^{-\frac{1}{2}} \cdot \|$ for any positive

self-adjoint A. Consider the following equation:

$$y = \mathcal{G}(u) + \eta. \tag{1}$$

Here $u \in X$ is an unknown function, $\mathcal{G}$ is the *forward response operator* mapping the unknown to the response space, $\eta \in Y$ is a noise and $y \in Y$ the observed data. We assume that η is a centered Gaussian with covariance Γ. The objective is to estimate u from y. We are particularly interested in situations where inversion of $\mathcal{G}$ on Y is ill-posed.

2.2. The Algorithms

We will show that the EnKF method for inverse problems, defined in detail in the next section, provides an estimate of u which lies in the subspace $\mathcal{A}$ spanned by the initial ensemble used to seed the algorithm. Via an iterative interacting particle system, initiated at this ensemble, the method provides solutions of the inverse problem which lie in the subspace $\mathcal{A}$ and are formed of combinations of the initial ensemble with nonlinear weights reflecting the observed data.

We wish to evaluate the EnKF method. Central to this task is the misfit functional

$$\Phi(u) = \|y - \mathcal{G}(u)\|_\Gamma^2, \tag{2}$$

which should be minimized in some sense. When inversion of $\mathcal{G}$ is ill-posed, minimization of Φ over the whole space X is not possible and regularization is required. For this paper three important regularization techniques [9] will play a role. The first is Tikhonov-Phillips regularization, whereby minimization of Φ is replaced by minimization of the regularized functional

$$I(u) = \|y - \mathcal{G}(u)\|_\Gamma^2 + \|(u - \bar{u})\|_C^2, \tag{3}$$

where C is a compact, positive and self-adjoint operator on X and $\bar{u}$ and element of X. The second is regularization via minimization of Φ over a compact set $\mathcal{A} \in X$. In this paper $\mathcal{A}$ will always be the linear span of a finite set of elements in X. The third is regularization via truncated iteration. The first, second and third methods can all be, and will be, used in conjunction. We note that minimization of I given by (3) has a statistical interpretation as the MAP estimator for Bayesian solution to the inverse problem with Gaussian prior $N(\bar{u}, C)$ on u [10]. Noting this intepretation, we observe that we will use as choices for $\mathcal{A}$ both a set of draws from the Gaussian prior $\mu_0 = N(\bar{u}, C)$ and a subset of the the eigenfunctions of C – the Karhunen-Loéve basis – added to $\bar{u}$. Finally we observe that, given the subspace $\mathcal{A}$, and given the truth $u^\dagger$ underlying the observed data y, we can compute the *best approximation* to the truth in $\mathcal{A}$. Since the EnKF method presented here finds solutions confined to $\mathcal{A}$ it is natural to compare it with the best approximation within $\mathcal{A}$. We now give more details of these least squares and best approximation algorithms used to benchmark the EnKF method. Furthermore, the following serves to establish a common notation used throughout our numerical studies.

Consider the space $\mathcal{A} = \text{span}\{\psi^{(j)}\}_{j=1}^{J}$ comprised of the initial ensemble used for the EnKF. Let μ_0 be some prior measure. The space $\mathcal{A}$ will be chosen based on the prior μ_0. When μ_0 is Gaussian $N(\bar{u}, C)$ let (λ_j, ϕ_j) denote eigenvalue/eigenvector pairs, in descending order by eigenvalue – this is the Karhunen-Loéve basis. However we do not restrict ourselves to Gaussian priors: for the Navier-Stokes example the prior will be the measure supported on the attractor. We will use: (a) $\psi^{(j)} \sim \mu_0$ i.i.d. and all algorithms using this choice of $\mathcal{A}$ will bear the subscript R for random; for Gaussian priors, we will additionally use (b) $\psi^{(j)} = \bar{u} + \sqrt{\lambda_j}\phi_j$, $j \leq J$, and all algorithms using this method will bear the subscript KL for Karhunen-Loéve .

We will compute EnKF_R and EnKF_KL approximations to the inverse problems of interest. For comparison we will compute the *best approximation*

$$u_{BA} = \text{argmin}_{u\in\mathcal{A}}||u - u^\dagger||_\Gamma^2. \tag{4}$$

leading to BA_R and BA_KL. Then, we will compute the projected basis *least squares* solution

$$u_{LS} = \text{argmin}_{u\in\mathcal{A}}||y - \mathcal{G}(u)||_\Gamma^2 \tag{5}$$

or generalizations to include Tikhonov-Phillip regularization [11], based on $\mu_0 = N(\bar{u}, C)$, and truncated Newton-CG iterative methods [12]. In all cases we will denote the resulting approximations denoted by LS_R and LS_KL.

2.3. Linear Problems

It is instructive to consider the case where $\mathcal{G}(u) = Gu$ for some linear operator $G : X \to Y$ as this will enable us to makes links with standard regularized least squares problems, the Kalman filter and hence the ensemble Kalman filter. If $\|y - G\cdot\|_\Gamma^2$ is continuous on $D(C)$ then the Tikhonov-Phillips regularized solution from (3) is given by $u_\text{TP} = \text{argmin}_{u\in D(C)} I(u)$, where

$$I(u) = \|y - Gu\|_\Gamma^2 + \|u - \bar{u}\|_C^2. \tag{6}$$

Furthermore this minimization has a unique solution which and may be computed explicitly to be

$$u_\text{TP} = \bar{u} + CG^*(GCG^* + \Gamma)^{-1}(y - G\bar{u}). \tag{7}$$

To see this we apply the Sherman-Morrison-Woodbury formula [13] to the normal equations resulting from minimization of I, as follows:

$$\begin{aligned}
u_\text{TP} &= \left(G^*\Gamma^{-1}G + C^{-1}\right)^{-1}\left(G^*\Gamma^{-1}y + C^{-1}\bar{u}\right) && (8)\\
&= \bar{u} + \left(G^*\Gamma^{-1}G + C^{-1}\right)^{-1}G^*\Gamma^{-1}(y - G\bar{u}) && (9)\\
&= \bar{u} + \left[C - CG^*(\Gamma + GCG^*)^{-1}GC\right]G^*\Gamma^{-1}(y - G\bar{u}) && (10)\\
&= \bar{u} + \left[CG^* - CG^*(\Gamma + GCG^*)^{-1}GCG^*\right]\Gamma^{-1}(y - G\bar{u}) && (11)\\
&= \bar{u} + \left[CG^* - CG^*\left(I - (\Gamma + GCG^*)^{-1}\Gamma\right)\right]\Gamma^{-1}(y - G\bar{u}) && (12)\\
&= \bar{u} + CG^*(\Gamma + GCG^*)^{-1}(y - G\bar{u}), && (13)
\end{aligned}$$

which we see immediately is exactly equation (7).

It is well-known that the Tikhonov regularized minimizer is the mean of the conditional random variable $u|y$ [14, 15], where the joint random variable is prescribed by the identities

$$u \sim N(\bar{u}, C) \tag{14}$$

$$y = Gu + \eta, \quad \eta \sim N(0, \Gamma) \perp u \tag{15}$$

and where $\perp$ denotes independence. Furthermore, we can connect this conditional random variable with the Kalman filter [16] as we now show. Let $Z = X \times Y$ and define $z \in Z$ and the linear operator $A : Z \to Z$ by

$$z = \begin{pmatrix} u \\ p \end{pmatrix}, \quad Az = \begin{pmatrix} u \\ Gu \end{pmatrix}.$$

Define the operator $H : Z \to Y$ by $H = (0, I)$, and $H^\perp : Z \to X$ by $H^\perp = (I, 0)$. Now consider the partially observed linear Gaussian dynamical system given by

$$z_{n+1} = Az_n,$$
$$y_{n+1} = Hz_{n+1} + \eta_{n+1}.$$

Here it is assumed that $\{\eta_n\}_{n\in\mathbb{Z}^+}$ is an i.i.d sequence with $\eta_1 \sim N(0, \Gamma)$. The Kalman filter computes the Gaussian random variable $z_n|\{y_m\}_{m=1}^n$, which is characterized by its mean and covariance for each n, and can be implemented sequentially as an iteration in n. If $u_0 \sim N(\bar{u}, C)$, $p_0 = Gu_0$ and $y_1 = y$ then, by construction, the first step of the Kalman filter $z_1|y_1$, projected into the u coordinate by $H^\perp$, is identical to the random variable $u|y$ defined previously. In particular the mean of one step of this Kalman filter is equal to u_{TP}. Note, further, that this link between the linear inverse problem suggests the intruiguing possibility of iterating the Kalman filter to get information about the inverse problem.

The EnKF proceeds by moving an ensemble of particles according to the mean Kalman update, and using this ensemble to approximate the covariance information needed to update the mean. In the limit of an infinite ensemble, EnKF converges to the Kalman Filter for linear problems. As we will show in section 3, when applied to a general class of nonlinear inverse problems, including the linear example under consideration here, the EnKF produces an approximation in the linear span of the initial ensemble. As such there is a natural connection to Krylov methods [17] and minimum residual methods such as MINRES [18] and GMRES [19] for appropriate choices of this initial ensemble.

Notice that we have made two interesting observations, in the context of this steady linear inverse problem: (i) we have connected it to a dynamic iterative estimation scheme; (ii) we have remarked that covariance information needed to implement the iteration can be approximated by an ensemble. These ideas can be generalized to nonlinear problems, and this is the subject of the next section.

3. EnKF Algorithm

The idea of the Kalman filter may be extended to nonlinear problems via the Extended Kalman Filter (ExKF) [20] which proceeds by (i) approximating the forward model by its linearization $\mathcal{G}(u) \approx \mathcal{G}(\bar{u}) + D\mathcal{G}(\bar{u})(u - \bar{u})$ and (ii) assuming that the desired distribution is well-approximated by a Gaussian $u \sim N(\bar{u}, C)$. It is then possible to apply the Kalman Filter, using this approximation. The EnKF [1] makes the further approximation of working with an ensemble and using this to estimate covariance information. For sequential problems numerical evidence demonstrates that the EnKF can produce reasonable state estimation even outside the regimes where the approximations (i) and (ii) are valid – see [21]. In this section we show how the EnKF can be used to construct an algorithm for inverse problems, by introducing an iterative scheme, generalizing subsection 2.3.

Define $Z = X \times Y$, $z \in Z$, $H : Z \to Y$ and $H^\perp : Z \to X$ as in the previous subsection. Then define $\Xi : Z \to Z$ by

$$\Xi(z) = \begin{pmatrix} u \\ \mathcal{G}(u) \end{pmatrix}.$$

We then estimate u by a variant of the EnKF algorithm applied to a partially observed dynamical system of the form

$$z_{n+1} = \Xi(z_n), \tag{16}$$

$$y_{n+1} = Hz_{n+1} + \eta_{n+1}. \tag{17}$$

Here it is assumed that $\{\eta_n\}_{n\in\mathbb{Z}^+}$ is an i.i.d. sequence with $\eta_1 \sim N(0, \Gamma)$. The EnKF estimates z_n given $\{y_m\}_{m=1}^n$. Since $H^\perp z = u$ we will apply the projector $H^\perp$ to our estimate of the state of the system at the n^{th} iterate of the algorithm in order to determine the unknown function u.

Recall that y is the given data. We will construct articifical data $\{y_n\}_{n\in\mathbb{Z}^+}$ from this single observation $y \in Y$ in two ways, both described below. The algorithm works with an ensemble of vectors $\{z_n^{(j)}\}_{j=1}^J$ which represent an ensemble of estimates of the state of the system at the n^{th} step of the algorithm, given the data $\{y_m\}_{m=1}^n$. We will take the mean of this ensemble as an estimate of the state. We assume that $z_0^{(j)} = \psi^{(j)}$ for $1 \le j \le J$. The set $\mathcal{A} = \text{span}\{\psi^{(j)}\}_{j=1}^J$ will play a key role in the analysis of the algorithm.

The algorithm breaks into two parts, a *prediction step* and an *analysis step*, and we desribe these separately. We then discuss properties of the resulting algorithm. The prediction step maps the current ensemble of particles into the data space, and thus introduces information about the forward model. The analysis step makes comparisons of the mapped ensemble, in the data space, with the data, or with noisily perturbed versions of the data; it is at this stage that the ensemble is modified in an attempt to better match the data.

3.1. Prediction Step

In this step we propagate the ensemble of particles forward under (16) giving

$$\widehat{z}_{n+1}^{(j)} = \Xi(z_n^{(j)}). \tag{18}$$

From this ensemble we define a sample mean and covariance as follows:

$$\overline{z}_{n+1} = \frac{1}{J}\sum_{j=1}^{J} \widehat{z}_{n+1}^{(j)} \tag{19}$$

$$C_{n+1} = \frac{1}{J}\sum_{j=1}^{J} \widehat{z}_{n+1}^{(j)}(\widehat{z}_{n+1}^{(j)})^T - \overline{z}_{n+1}\overline{z}_{n+1}^T. \tag{20}$$

All the vectors and operators involved have block structure inherited from the structure of the space $Z = X \times Y$. For example we have

$$\widehat{z}_{n+1}^{(j)} = \begin{pmatrix} \widehat{u}_{n+1}^{(j)} \\ \widehat{p}_{n+1}^{(j)} \end{pmatrix} = \begin{pmatrix} u_n^{(j)} \\ \mathcal{G}(u_n^{(j)}) \end{pmatrix}, \quad z_n^{(j)} = \begin{pmatrix} u_n^{(j)} \\ p_n^{(j)} \end{pmatrix}.$$

We also have

$$\overline{z}_n = \begin{pmatrix} \overline{u}_n \\ \overline{p}_n \end{pmatrix}, \quad C_n = \begin{pmatrix} C_n^{uu} & C_n^{up} \\ (C_n^{up})^T & C_n^{pp} \end{pmatrix}.$$

The vectors $\overline{u}_n$ and $\overline{p}_n$ are given by

$$\overline{u}_n = \frac{1}{J}\sum_{j=1}^{J} \widehat{u}_n^{(j)} = \frac{1}{J}\sum_{j=1}^{J} u_{n-1}^{(j)} \tag{21}$$

$$\overline{p}_n = \frac{1}{J}\sum_{j=1}^{J} \widehat{p}_n^{(j)} = \frac{1}{J}\sum_{j=1}^{J} \mathcal{G}(u_{n-1}^{(j)}) \tag{22}$$

The blocks within C_n are given by

$$C_n^{uu} = \frac{1}{J}\sum_{j=1}^{J} \widehat{u}_n^{(j)}(\widehat{u}_n^{(j)})^T - \overline{u}_n\overline{u}_n^T, \tag{23}$$

$$C_n^{up} = \frac{1}{J}\sum_{j=1}^{J} \widehat{u}^{(j)}(\widehat{p}^{(j)})^T - \overline{u}_n\overline{p}_n^T, \tag{24}$$

$$C_n^{pp} = \frac{1}{J}\sum_{j=1}^{J} \widehat{p}^{(j)}(\widehat{p}^{(j)})^T - \overline{p}_n\overline{p}_n^T, . \tag{25}$$

3.2. Analysis Step

The analysis step combines the ensemble $\{\widehat{z}_{n+1}^{(j)}\}_{j=1}^J$ with data $\{y_{n+1}^{(j)}\}_{j=1}^J$, using ideas stemming from the Kalman filter, but with the predicted mean and covariance estimated from the samples. For the *perturbed observation* variant of EnKF the noisly perturbed data is constructed as follows:

$$y_{n+1}^{(j)} = y + \eta_{n+1}^{(j)}. \tag{26}$$

Here the $\eta_{n+1}^{(j)}$ are an i.i.d collections of vectors indexed by (j,n) with $\eta_1^{(1)} \sim N(0,\Gamma)$. We will also study, for comparison, an unperturbed algorithm where we simply take

$$y_{n+1}^{(j)} = y. \tag{27}$$

We do not include the results for this case, as they are sufficiently similar to the results for the case of perturbed observations.

To construct the analysis step we define the operators S_n and K_n by

$$S_n = HC_nH^T + \Gamma, \tag{28}$$

$$K_n = C_nH^TS_n^{-1}, \tag{29}$$

where

$$H^T = \begin{pmatrix} 0 \\ I \end{pmatrix}.$$

The operator K_{n+1} is the Kalman gain and the particles are updated according to the standard Kalman update formula which blends model and data:

$$z_{n+1}^{(j)} = I\widehat{z}_{n+1}^{(j)} + K_{n+1}(y_{n+1}^{(j)} - H\widehat{z}_{n+1}^{(j)}) \tag{30}$$

$$= (I - K_{n+1}H)\widehat{z}_{n+1}^{(j)} + K_{n+1}y_{n+1}^{(j)}. \tag{31}$$

3.3. Properties of the Algorithm

Recall the space $\mathcal{A}$ formed from the linear span of the initial ensemble. All subsequent EnKF vectors remain within this space.

Theorem 3.1. *For every $(n,j) \in \mathbb{N} \times \{1,\cdots,J\}$ we have $z_n^{(j)} \in \mathcal{A}$.*

Proof. This is a straightforward induction, using the properties of the update formulae. Clearly the statement is true for $n = 0$. Assume that it is true for n. The operators S_n and K_n have particular structure inherited from the form of H. In particular we have

$$S_n = C_n^{pp} + \Gamma$$

and

$$K_n = \begin{pmatrix} C_n^{up}(C_n^{pp} + \Gamma)^{-1} \\ C_n^{pp}(C_n^{pp} + \Gamma)^{-1} \end{pmatrix}.$$

Note that

$$C_n^{up} = \frac{1}{J}\sum_{j=1}^{J} \widehat{u}_n^{(j)}(\widetilde{p}_n^{(j)})^T$$

and

$$C_n^{pp} = \frac{1}{J}\sum_{j=1}^{J} \widehat{p}_n^{(j)}(\widetilde{p}_n^{(j)})^T$$

where

$$\widetilde{p}_n^{(j)} = \widehat{p}_n^{(j)} - \frac{1}{J}\sum_{\ell=1}^{J} \widehat{p}_n^{(\ell)}.$$

Recall that, from the structure of the map Ξ, we have

$$\widehat{u}_{n+1}^{(j)} = u_n^{(j)}, \quad \widehat{p}_{n+1}^{(j)} = \mathcal{G}(u_n^{(j)})$$

and from the definition of H we get that

$$K_n H = \begin{pmatrix} 0 & C_n^{up}(C_n^{pp} + \Gamma)^{-1} \\ 0 & C_n^{pp}(C_n^{pp} + \Gamma)^{-1} \end{pmatrix}$$

Using these facts in (30) we deduce that the update equations are

$$u_{n+1}^{(j)} = u_n^{(j)} + C_{n+1}^{up}(C_{n+1}^{pp} + \Gamma)^{-1}\Big(y_{n+1}^{(j)} - \mathcal{G}(u_n^{(j)})\Big) \tag{32}$$

$$p_{n+1}^{(j)} = \mathcal{G}(u_n^{(j)}) + C_{n+1}^{pp}(C_{n+1}^{pp} + \Gamma)^{-1}\Big(y_{n+1}^{(j)} - \mathcal{G}(u_n^{(j)})\Big). \tag{33}$$

If we define

$$d_{n+1}^{(j)} = (C_{n+1}^{pp} + \Gamma)^{-1}\Big(y_{n+1}^{(j)} - \mathcal{G}(u_n^{(j)})\Big)$$

then the update formula (32) for the unknown u may be written as

$$u_{n+1}^{(j)} = u_n^{(j)} + \frac{1}{J}\sum_{k=1}^{J} \langle \widetilde{p}_{n+1}^{(k)}, d_{n+1}^{(j)} \rangle \widehat{u}_{n+1}^{(k)} \tag{34}$$

$$= u_n^{(j)} + \frac{1}{J}\sum_{k=1}^{J} \langle \widetilde{p}_{n+1}^{(k)}, d_{n+1}^{(j)} \rangle u_n^{(k)}. \tag{35}$$

In view of the inductive hypothesis at step n, this demonstrates that $u_{n+1}^{(j)} \in \mathcal{A}$ for all $j \in \{1, \cdots, J\}$ and the proof is complete. □

Remark 3.2. *The proof demonstrates that the solution at step n is simply a linear combination of the original samples; however the coefficients in the linear combination depend nonlinearly on the process. Note also that the update formula (33) for the system response may be written as*

$$p_{n+1}^{(j)} = \mathcal{G}(u_n^{(j)}) + \frac{1}{J}\sum_{k=1}^{J} \langle \widetilde{p}_{n+1}^{(k)}, d_{n+1}^{(j)} \rangle \widehat{p}_{n+1}^{(k)} \tag{36}$$

$$= \mathcal{G}(u_n^{(j)}) + \frac{1}{J}\sum_{k=1}^{J} \langle \widetilde{p}_{n+1}^{(k)}, d_{n+1}^{(j)} \rangle \mathcal{G}(u_n^{(k)}). \tag{37}$$

Let $u^\dagger$ denote the true unknown function which underlies the data so that

$$y = \mathcal{G}(u^\dagger) + \eta^\dagger$$

for some noise $\eta^\dagger \in Y$. We have the following lower bound for the accuracy of the estimates produced by the EnKf algorithm for inverse problems:

Corollary 3.3. *The error between the estimate of u at time n and the truth $u^\dagger$ satisfies*

$$\|\overline{u}_n - u^\dagger\| \geq \inf_{v \in \mathcal{A}} \|v - u^\dagger\|.$$

4. Elliptic Equation

As a simple pedagogical example, we consider the ill-posed inverse problem of recovering the right-hand side of an elliptic equation in one spatial dimension given noisy observation of the solution. This explicitly solvable linear model will allow us to elucidate the performance of EnKF as an iterative regularisation method. The results show that EnKF performs comparably to Tikhonov-Phillips regularized least squares for this problem. The lower bound BA is always significantly better, but of course unimplementable in practice, since the truth is unknown.

4.1. Setting

Consider the one dimensional elliptic equation

$$\begin{aligned} -\frac{d^2p}{dx^2} + p &= u \\ u(0) = u(\pi) &= 0. \end{aligned}$$

Thus $G = A^{-1}$ where $A = (-\frac{d^2}{dx^2} + 1)$ and $D(A) = H^2(I) \cap H_0^1(I)$ with $I = (0, \pi)$. We are interested in the inverse problem of recovering u from noisy observations of p:

$$\begin{aligned} y &= p + \eta \\ &= A^{-1}u + \eta. \end{aligned}$$

For simplicity we assume that the noise is white: $\eta \sim N(0, \gamma^2 I)$. We consider Tikhonov-Phillips regularization of the form $\|u\|_C^2$, where $C = \beta(A - I)^{-1}$. The problem may be solved explicitly in the Fourier sine basis, and the coefficients of the Tikhonov-Phillips regularized least squares solution, in this basis, u_k may be expressed in terms of the coefficients of the data in the same basis, y_k:

$$\left[\left(\frac{1}{\gamma(1+k^2)}\right)^2 + \beta^{-1}k^2\right] u_k = \frac{y_k}{\gamma^2(1+k^2)}, \quad k = 1, ..., \infty \tag{38}$$

This demonstrates explicitly the regularization which is present for wavenumbers k such that $k^6 \geq \mathcal{O}(\beta\gamma^{-2})$. We now present some detailed numerical experiments which will place the EnKF algorithm in the context of an iterative regularization scheme. This will provide a roadmap for understanding the nonlinear inverse problem applications we explore in subsequent sections.

4.2. Numerical Results

Throughout this subsection we choose $\beta = 10$ and $\gamma = 0.01$. We choose a truth $u^\dagger \sim N(0, C)$ and simulate data from the model: $y = Gu^\dagger + \eta^\dagger$, where $\eta^\dagger \sim \mathcal{N}(0, \Gamma)$. Recall from Section 2 that the space $\mathcal{A} = \text{span}\{\psi^{(j)}\}_{j=1}^J$ will be chosen either based on draws from $N(0, C)$ (with subscript R) or from the Karhunen-Loéve basis for $N(0, C)$ (with subscript KL.)

We evaluate the lower bounds on the error resulting from best approximation in the random and Karhunen-Loéve bases respectively: BA_R and BA_{KL}. We also compute errors in the projected basis regularized least squares solutions LS_R and LS_{KL} with additional Tikhonov-Phillips regularisation to compensate for the large number of basis vectors. And finally, we compute the errors resulting from the EnKF means, namely $EnKF_R$ and $EnKF_{KL}$, as detailed in the previous section. Notice that for the EnKF methods we may take just one-step, or we may iterate over many steps; experiments describing both are reported.

In Figure 1 (top) we show three examples of the error of the estimators with respect to the truth over 100 different ensembles $\mathcal{A}$ of random draws from the prior: one step of $\mathrm{EnKF_R}$, Tikhonov-Phillips regularized $\mathrm{LS_R}$ (equivalent to KF) and $\mathrm{BA_R}$. The methods clearly indicate that the EnKF method is comparable to the LS method in terms of accuracy, but involves no matrix inversions. Both EnKF and LS are less accurate than BA, of course, but produce errors of similar order of magnitude.

In Figure 1 (bottom) the results for random draws, with suffix R, are averaged over many initial choices of ensembles, and compared with results using the Karhunen-Loéve basis, with suffix KL. The errors with suffix R are therefore averaged over randomly drawn initial bases and horizontal lines indicate these quantities in the plot. EnKF based methods are compared with BA and LS type methods. These results again show that EnKF and LS are similar in terms of accuracy, and produce errors of similar order of magnitude to BA. Furthermore they show that, when iterated, the error of $\mathrm{EnKF_{KL}}$ decreases for some time before reaching its minimum, while EnKF_R reaches its minimum after a small number of iterations (often only one), and then increases (note that the averaging of the error hides the actual value at which the error starts increasing for any given choice of initial ensemble).

5. Groundwater flow

Next, we will investigate an inverse problem arising in reservoir modeling. Typically, the calibration of subsurface flow models consist of estimating geologic properties whose model predictions "best fit" the flow measurements [22]. In other words, the calibration can be posed as the inverse problem of finding the geologic properties given subsurface flow data. In particular, herein we will consider the estimation of the conductivity of an aquifer form hydraulic head measurements. Similarly to the previous section, we find comparable performance of EnKF and LS, with the latter here regularized in a Newton-CG fashion. Again, BA is included for comparison.

5.1. Setting

We consider groundwater flow in a two-dimensional confined aquifer whose physical domain is $\Omega = [0, 6] \times [0, 6]$. The hydraulic conductivity is denoted by K. The flow in the aquifer, is described in terms of the piezometric head $h(x)$ ($x \in \Omega$) which, in the

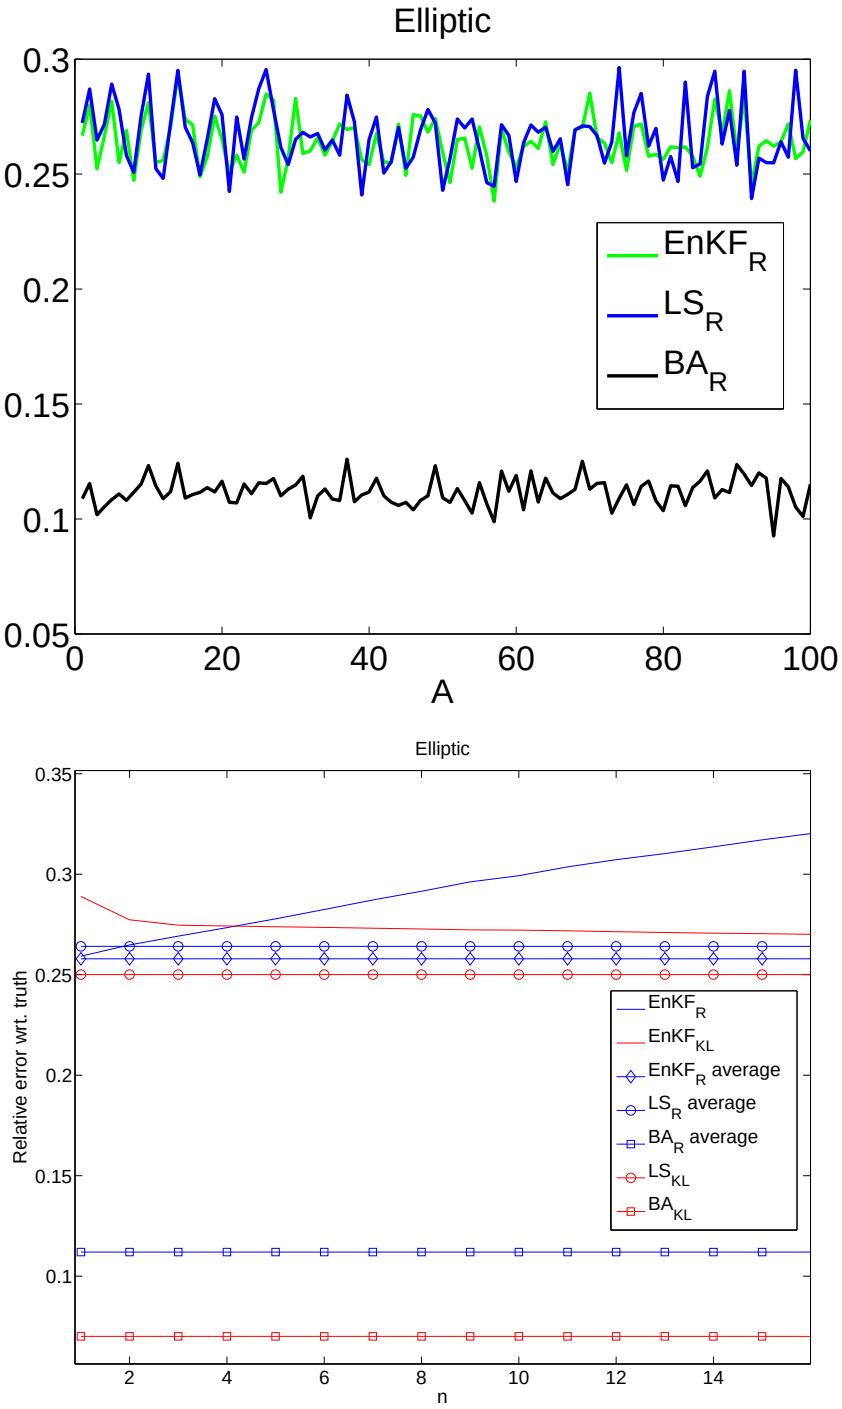

Figure 1. Top: Comparison over an ensemble of random ensembles $\mathcal{A}$ of the relative errors of one iteration of $\mathrm{EnKF_R}$ versus (Tikhonov regularized) least squares ($\mathrm{LS_R}$) and the best approximation ($\mathrm{BA_R}$). Bottom: Comparison of one trajectory of $\mathrm{EnKF_R}$ and $\mathrm{EnKF_{KL}}$ over iteration n, as well as $\mathrm{EnKF_R}$, $\mathrm{BA_R}$ and $\mathrm{LS_R}$ averaged over ensembles, and $\mathrm{BA_{KL}}$ and $\mathrm{LS_{KL}}$.

steady-state is governed by the following equation [23]

$$-\nabla \cdot e^u \nabla h = f \qquad \text{in} \quad \Omega \tag{39}$$

where $u \equiv \log K$ and f is defined by

$$f(x_1, x_2) = \begin{cases} 0 & \text{if } 0 < x_2 \leq 4, \\ 137 & \text{if } 4 < x_2 < 5, \\ 274 & \text{if } 5 \leq x_2 < 6. \end{cases} \tag{40}$$

We consider the following boundary conditions

$$h(x, 0) = 100, \qquad \frac{\partial h}{\partial x}(6, y) = 0, \qquad -e^u \frac{\partial h}{\partial x}(0, y) = 500, \qquad \frac{\partial h}{\partial y}(x, 6) = 0, \tag{41}$$

For the physical interpretation of the source term and boundary conditions (40)-(41), we refer the reader to [24] where a similar model was used as a benchmark for inverse modeling in groundwater flow. A similar model was also studied in [25, 26] also in the context of parameter identification. We will be interested in the inverse problem of recovering the hydraulic conductivity, or more precisely its logarithm u, from noisy pointwise measurements of the piezometric head h. This is a model for the situation in groundwater applications where observations of head are used to infer the conductivity of the aquifer.

5.2. Numerical Results

We now explain in detail the numerical results presented in Figures 2 and 3. We let $\mathcal{G}(u) = \{p(x_k)\}_{k \in \mathbb{K}}$ where $\mathbb{K}$ is some finite set of points in Ω with cardinality K. In particular, $\mathbb{K}$ is given by the configuration of $N = 100$ observation wells displayed in Figure 3 (Top right). We introduce the prior Gaussian $\mu_0 \sim \mathcal{N}(\overline{u}, C)$, let $u^\dagger \sim \mu_0$, and simulate data $y = \mathcal{G}(u^\dagger) + \eta^\dagger$, where $\eta^\dagger \sim \mathcal{N}(0, \Gamma)$. As in the previous section, we let $C = \beta(-\Delta)^{-\alpha}$ and $\Gamma = \gamma^2 I$, and here we choose $\alpha = 1.3, \beta = 0.5, \overline{u} = 4$ and $\gamma = 7$ fixed. We reiterate from Section 2 that the space $\mathcal{A} = \text{span}\{\psi^{(j)}\}_{j=1}^J$ will be chosen based on either draws from the prior μ_0, with subscript R, or on the Karhunen-Loéve basis, with subscript KL.

The forward model (39)-(41) is discretized with cell-centered finite differences [27]. For the approximation of the LS problem, we implemented the Newton-CG method of [12]. We conduct experiments analogous to those of the previous section and the results are shown in Figure 2 and Figure 3. The results in Figure 2 are very similar to those shown in the previous section for the linear elliptiuc problem, demonstrating the robustness of the observations made in that section for the solution of inverse problems in general, using the EnKF methodology herein.

6. Navier-Stokes Equation

In this section, we consider an inverse problem in fluid dynamics, which is relevant to data assimilation applications in oceanography and meteorology. In particular, we examine the problem of recovering the initial condition of the Navier-Stokes Equation (NSE), given noisy pointwise observations of the velocity field at later times. We will investigate a regime in which the combination of viscosity, time-interval, and truncation

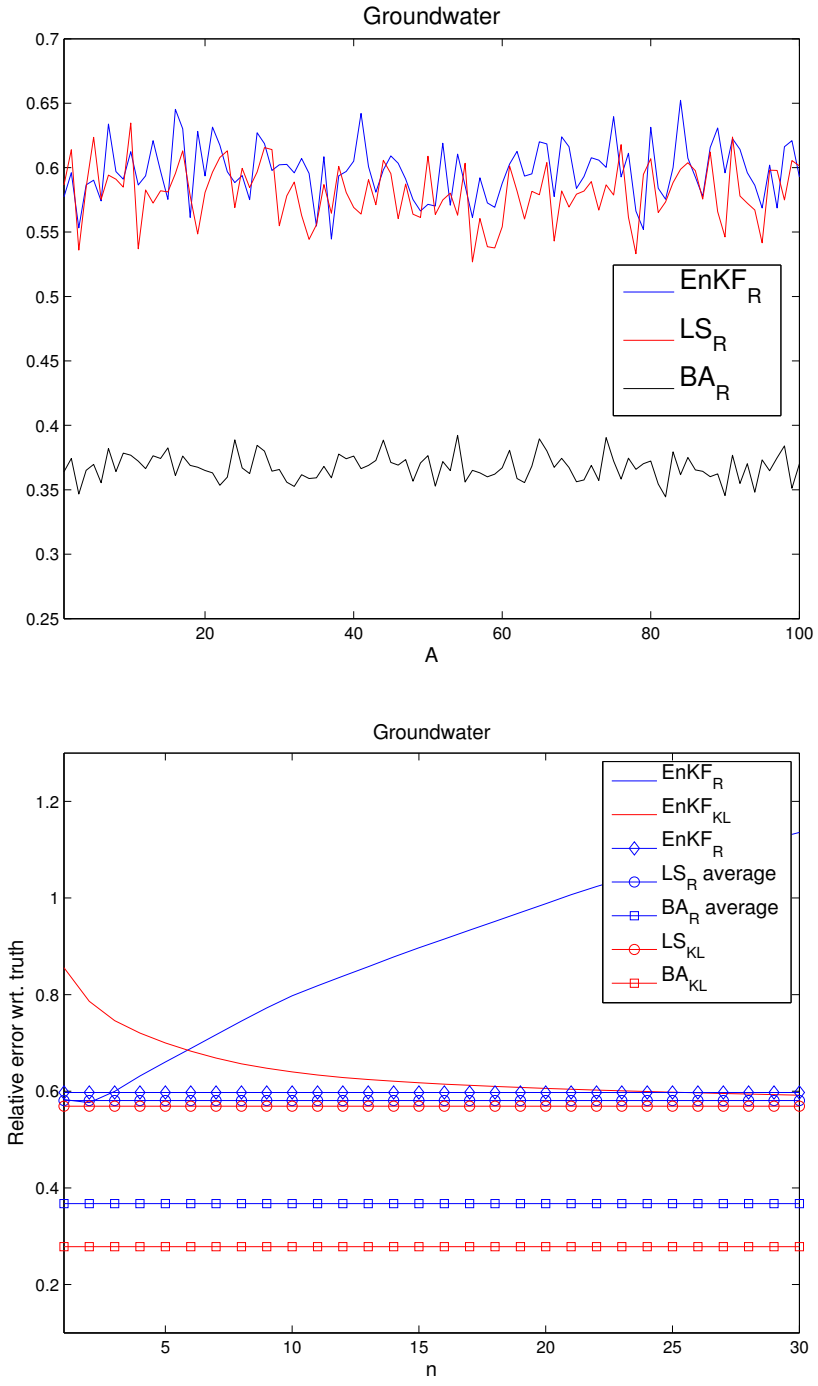

Figure 2. Top: Comparison over an ensemble of random ensembles $\mathcal{A}$ of the relative errors of one iteration of $\mathrm{EnKF_R}$ versus least squares ($\mathrm{LS_R}$) and the best approximation ($\mathrm{BA_R}$). Bottom: Comparison of one trajectory of $\mathrm{EnKF_R}$ and $\mathrm{EnKF_{KL}}$ over iteration n, as well as $\mathrm{EnKF_R}$, $\mathrm{BA_R}$ and $\mathrm{LS_R}$ averaged over ensembles, and $\mathrm{BA_{KL}}$ and $\mathrm{LS_{KL}}$.

of the forward model is such that the inverse problem is only "mildly ill-posed", in the sense that it does not exhibit the ill-posedness of the function space limit.

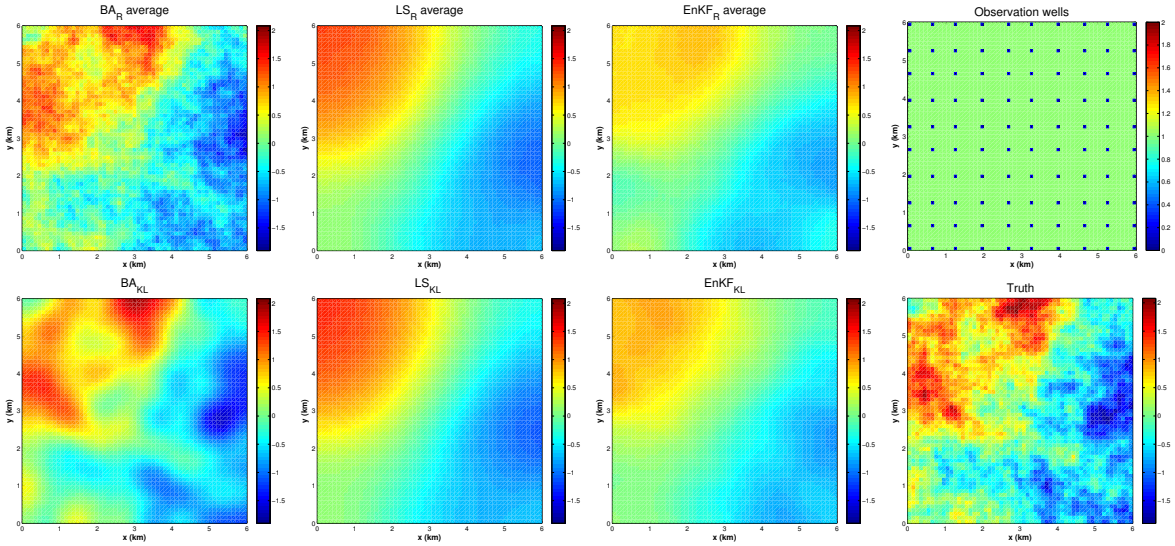

Figure 3. Estimated log K. Top left: $\mathrm{BA_R}$, average. Top middle left: $\mathrm{LS_R}$, average. Top middle right: first iteration of $\mathrm{EnKF_R}$, average. Top right: measurement well locations. Bottom left: $\mathrm{BA_{KL}}$. Bottom middle left: $\mathrm{LS_{KL}}$. Bottom middle right: $\mathrm{EnKF_{KL}}$. Bottom right: the truth $u^\dagger$.

In practice, a regularization term is included anyway, but with the purpose of incorporating background, or climatological, information into the model in the form of a prior.

As in the previous two sections we will see that EnKF and LS type methods perform comparably, in both the case of random and Karhunen-Loéve based initial draws. Furthermore the lower bound produced by BA type methods is of similar order of magnitude to the EnKF-based methods although the actual error is, of course, smaller. However, we will also see that the behavior of the iterated $\mathrm{EnKF_R}$ is quite different from what we observed in the previous sections. We conjecture that this is because of the mildly ill-posed nature of this problem. Indeed we have have repeated the results of this section at higher viscosity, where linear damping in the forward model induces greater ill-posedness, and confirmed that we recover results for the iterated $\mathrm{EnKF_R}$ which are similar to those in the previous sections.

6.1. Setting

We consider the 2D Navier-Stokes equation on the torus $\mathbb{T}^2 := [-1, 1) \times [-1, 1)$ with periodic boundary conditions:

$$\begin{aligned} \partial_t v - \nu \Delta v + v \cdot \nabla v + \nabla p &= f \quad \text{for all } (x,t) \in \mathbb{T}^2 \times (0,\infty), \\ \nabla \cdot v &= 0 \quad \text{for all } (x,t) \in \mathbb{T}^2 \times (0,\infty), \\ v &= u \quad \text{for all } (x,t) \in \mathbb{T}^2 \times \{0\}. \end{aligned}$$

Here $v\colon \mathbb{T}^2 \times (0,\infty) \to \mathbb{R}^2$ is a time-dependent vector field representing the velocity, $p\colon \mathbb{T}^2 \times (0,\infty) \to \mathbb{R}$ is a time-dependent scalar field representing the pressure, $f\colon \mathbb{T}^2 \to \mathbb{R}^2$ is a vector field representing the forcing (which we assume to be time-independent

for simplicity), and ν is the viscosity. We are interested in the inverse problem of determining the initial velocity field u from pointwise measurements of the velocity field at later times. This is a model for the situation in weather forecasting where observations of the atmosphere are used to improve the initial condition used for forecasting.

6.2. Numerical Results

We let $t_j = jh$, for $j = 0, \ldots, J$, and $\mathcal{G}(u) = \{v(x_k, t_j)\}_{(j,k)\in\mathbb{K}'}$ where $\mathbb{K}' = \{0, \cdots, J\} \times \mathbb{K}$ and $\mathbb{K}$ is some finite set of points in Ω with cardinality K. In particular, we take $\mathbb{K}$ to be the set of grid points in physical space implied by the underlying spectral truncation used in the numerical integration (details below). As discussed in Section 2, we introduce a prior μ_0, let $u^\dagger \sim \mu_0$, and simulate data $y = \mathcal{G}(u^\dagger) + \eta^\dagger$, where $\eta^\dagger \sim \mathcal{N}(0, \Gamma)$. As in the previous section, we let $\Gamma = \gamma^2 I$, so our tunable parameter is γ which we fix at $\gamma = 0.01$ for these experiments. The prior μ_0 is defined to be the empirical measure supported on the attractor, i.e. it is defined by samples of a trajectory of the forward model after convergence to statistical equilibrium. We use 10^4 time-steps to construct this empirical measure. Once again the space $\mathcal{A} = \text{span}\{\psi^{(j)}\}_{j=1}^J$ is chosen based on J samples from the (now non-Gaussian) prior μ_0, or on a Karhunen-Loéve expansion based on the empirical mean $\bar{u}$ and covariance C from the long forward simulation giving rise to μ_0. It is important to note that the truth $u^\dagger$ is included in the long trajectory used to construct μ_0. Therefore, some of the initial ensembles drawn from μ_0 end up containing a snapshot very close to the truth; such results are overly optimistic as this is not the typical situation. It is interesting that each method has comparably overly optimistic results for such ensembles. For (ii), we use the the empirical mean $\bar{u}$ and covariance C over a long forward simulation.

The forcing in f is taken to be $f = \nabla^\perp \psi$, where $\psi = \cos(\pi k \cdot x)$ and $\nabla^\perp = J\nabla$ with J the canonical skew-symmetric matrix, and $k = (5, 5)$. The method used to approximate the forward model is a modification of a fourth-order Runge-Kutta method, ETD4RK [28], in which the Stokes semi-group is computed exactly by working in the incompressible Fourier basis $\{\psi_k(x)\}_{k\in\mathbb{Z}^2\setminus\{0\}}$, and Duhamel's principle (variation of constants formula) is used to incorporate the nonlinear term. We use a time-step of $dt = 0.005$. Spatially, a Galerkin spectral method [29] is used, in the same basis, and the convolutions arising from products in the nonlinear term are computed via FFTs. We use a double-sized domain in each dimension, padded with zeros, resulting in 64^2 grid-point FFTs, and only half the modes in each direction are retained when transforming back into spectral space again. This prevents aliasing, which is avoided as long as more than one third of the domain in which FFTs are computed consists of such padding with zeros. The dimension of the attractor is determined by the viscosity parameter ν. For the particular forcing used there is an explicit steady state for all $\nu > 0$ and for $\nu \geq 0.035$ this solution is stable (see [30], Chapter 2 for details). As ν decreases the flow becomes increasingly complex and we focus subsequent studies of the inverse problem on the mildly chaotic regime which arises for $\nu = 0.01$. Regarding observations, we

let $h = 4 \times dt = 0.02$ and take $J = 10$, so that $T = 0.2$. The observations are made at all numerically resolved, and hence observable, wavenumbers in the system; hence $K = 32^2$, because of the padding to avoid aliasing.

The numerical results resulting from these experiments are displayed in in Figures 4 and 5. The results are very similar to those of the previous two sections, qualitatively: EnKF and LS type methods perform comparably, in both the case of random and Karhunen-Loéve based initial draws; furthermore the lower bound produced by BA type methods is of similar order of magnitude to the EnKF-based methods although the actual error is, of course, smaller. However, we also see that the behavior of the iterated EnKF_R is quite different from what we observed in the previous sections since it decreases montonically. We conjecture that this is because of the mildly ill-posed nature of this problem. We also note that Tikhonov of truncated iterative regularisation is not necessary for *LS* for this example, because of the truncation of wavenumbers, the low viscosity, and the short time-interval. We have have repeated the results of this section at the higher viscosity $\nu = 0.1$, where linear damping in the forward model induces greater ill-posedness, and confirmed that we recover results for the iterated EnKF_R which are similar to those in the previous sections, in that the error displays a minimum after one or two steps Also, for this higher value of ν, Tikhonov regularization is used to ensure that LS delivers reasonable results, because of the greater ill-posedness. Finally, we have checked that the behavior of the error in the iterated EnKF_R is repeatable for Gaussian prior μ_0, and hence is not a result of non-Gaussian ensembles used in the figures.

7. Conclusions

We have illustrated the use of EnKF as a derivative-free optimization tool for inverse problems, showing that the method computes a nonlinear approximation in the linear span of the initial ensemble. We have also demonstrated comparable accuracy to least squares based methods and shown that, furthermore, the accuracy is of the same order of magnitude as best approximation within the linear span of the initial ensemble. Further study of the EnKF methodology for inverse problems, and in particular its accuracy with respect to choice of initial ensemble, would be of interest. Furthermore, in this paper we have concentrated purely on the accuracy of state estimation using EnKF. Study of its accuracy in terms of uncertainty quantification will yield further insight. Finally, although we have studied a time-dependent example (the Navier-Stokes equation) we did not use a methodology which exploited the sequential aquisition of the data – we concatenated all the data in space and time. In future work it is of interest to study ideas similar to those herein, but exploiting sequential structure.

Acknowledgments

This was supported by the ERC, EPSRC & ONR.

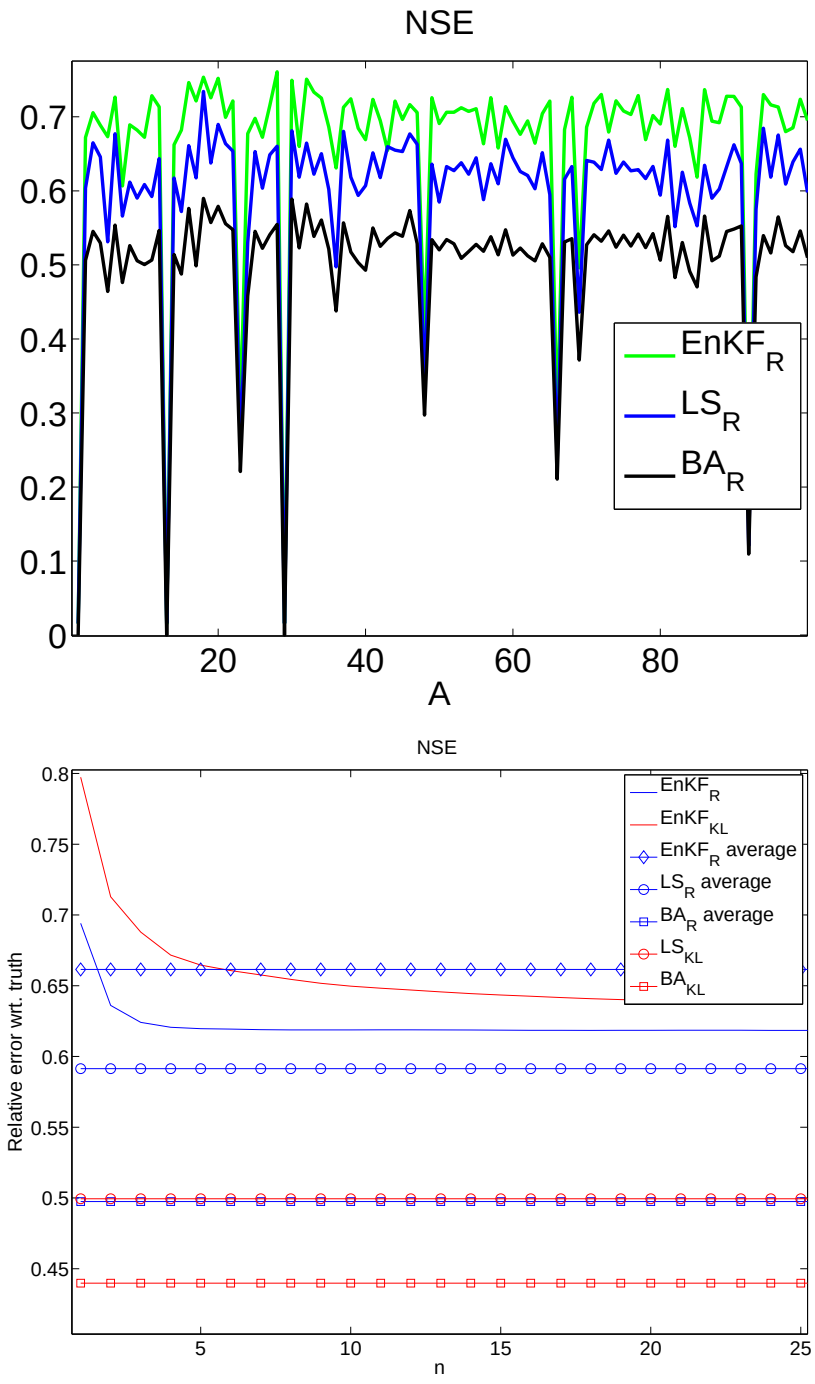

Figure 4. Top: Comparison over an ensemble of random ensembles $\mathcal{A}$ of the relative errors of one iteration of $\mathrm{EnKF_R}$ versus least squares ($\mathrm{LS_R}$) and the best approximation ($\mathrm{BA_R}$). Bottom: Comparison of one trajectory of $\mathrm{EnKF_R}$ and $\mathrm{EnKF_{KL}}$ over iteration n, as well as $\mathrm{EnKF_R}$, $\mathrm{BA_R}$ and $\mathrm{LS_R}$ averaged over ensembles, and $\mathrm{BA_{KL}}$ and $\mathrm{LS_{KL}}$.

References

[1] G. Evensen. Sequential data assimilation with a nonlinear quasi-geostrophic model using monte carlo methods to forecast error statistics. *J. Geophysical Research-All Series-*, 99:10–10, 1994.

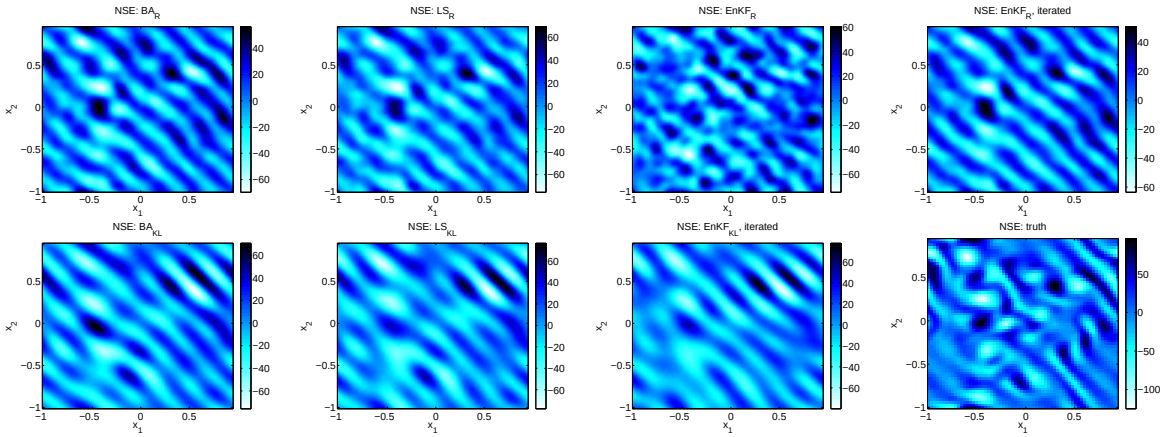

Figure 5. Vorticity of initial condition $\omega(0)$. Top left: $\mathrm{BA_R}$, average. Top middle left: $\mathrm{LS_R}$, average. Top middle right: first iteration of $\mathrm{EnKF_R}$, average. Top right: tenth iteration of $\mathrm{EnKF_R}$, average. Bottom left: $\mathrm{BA_{KL}}$. Bottom middle left: $\mathrm{LS_{KL}}$. Bottom middle right: $\mathrm{EnKF_{KL}}$. Bottom right: the truth $u^\dagger$.

[2] G. Evensen and P.J. Van Leeuwen. Assimilation of geosat altimeter data for the agulhas current using the ensemble kalman filter with a quasi-geostrophic model. *Monthly Weather*, 128:85–96, 1996.

[3] N. Geir, L. Johnsen, S. Aanonsen, and E. Vefring. Reservoir monitoring and continuous model updating using ensemble kalman filter. In *SPE Annual Technical Conference and Exhibition*, 2003.

[4] P.L. Houtekamer and H.L. Mitchell. A sequential ensemble kalman filter for atmospheric data assimilation. *Monthly Weather Review*, 129(1):123–137, 2001.

[5] G. Evensen. *Data assimilation: the ensemble Kalman filter*. Springer Verlag, 2009.

[6] E. Kalnay. *Atmospheric modeling, data assimilation, and predictability*. Cambridge Univ Pr, 2003.

[7] D.S. Oliver, A. C. Reynolds, and N. Liu. *Inverse Theory for Petroleum Reservoir Characterization and History Matching*. Cambridge University Press, ISBN: 9780521881517, 1st edition, 2008.

[8] G. Li and A. Reynolds. An iterative ensemble kalman filter for data assimilation. In *SPE Annual Technical Conference and Exhibition*, 2007.

[9] H.W. Engl, M. Hanke, and A. Neubauer. *Regularization of inverse problems*, volume 375. Springer, 1996.

[10] J.P. Kaipio and E. Somersalo. *Statistical and computational inverse problems*. Springer Science+ Business Media, Inc., 2005.

[11] C.R. Vogel. *Computational methods for inverse problems*. Society for Industrial Mathematics, 2002.

[12] M. Hanke. Regularizing properties of a truncated newton-cg algorithm for nonlinear inverse problems. *Numerical Functional Analysis and Optimization*, 18(9-10):971–993, 1997.

[13] N.J. Higham. *Accuracy and stability of numerical algorithms*. Number 48. Siam, 1996.

[14] A. Mandelbaum. Linear estimators and measurable linear transformations on a hilbert space. *Prob. Theo. Rel. Fields*, 65:385–397, 1984.

[15] M.S. Lehtinen, L. Paivarinta, and E. Somersalo. Linear inverse problems for generalized random variables. *Inverse Problems*, 5:599–612, 1989.

[16] R.E. Kalman. A new approach to linear filtering and prediction problems. *Journal of basic Engineering*, 82(Series D):35–45, 1960.

[17] Y. Saad. *Iterative methods for sparse linear systems*. PWS Pub. Co., 1996.

[18] C.C. Paige and M.A. Saunders. Solution of sparse indefinite systems of linear equations. *SIAM*

Journal on Numerical Analysis, pages 617–629, 1975.

[19] Y. Saad and M.H. Schultz. Gmres: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM J. Sci. Stat. Comput.*, 7(3):856–869, 1986.

[20] A.H. Jazwinski. *Stochastic processes and filtering theory*. Academic Pr, 1970.

[21] K.J.H. Law and A.M. Stuart. Evaluating data assimilation algorithms. *Accepted for publication in Monthly Weather Review. Arxiv preprint arXiv:1107.4118*, 2011.

[22] S.I. Aanonsen, G. Naevdal, D.S. Oliver, A.C. Reynolds, , and B Valles. The Ensemble Kalman Filter in Reservoir Engineering–a Review. *SPE J.*, 14(3):393–412, 2009.

[23] J. Bear. *Dynamics of Fluids in Porous Media*. Dover Pulications, New York, 1972.

[24] J. Carrera and S. P. Neuman. Estimation of aquifer parameters under transient and steady state conditions: 3. application to synthetic and field data. *Water Resources Research*, 22.

[25] M. Hanke. A regularizing Levenberg-Marquardt scheme, with applications to inverse groundwater filtration problems. *Inverse Problems*, 13:79–95, 1997.

[26] M. A. Iglesias and C. Dawson. The respresenter method for state and parameter estimation in single-phase Darcy flow. *Computer Methods in Applied Mechanics and Engineering*, 196:45774596, 2007.

[27] T.F. Russell and M.F. Wheeler. Finite element and finite difference methods for continuous flows in porous media. *In: R.E. Ewing, Editor, Mathematics of Reservoir Simulation*, SIAM,Philadelphia, PA.

[28] SM Cox and PC Matthews. Exponential time differencing for stiff systems. *Journal of Computational Physics*, 176(2):430–455, 2002.

[29] J.S. Hesthaven, S. Gottlieb, and D. Gottlieb. *Spectral methods for time-dependent problems*, volume 21. Cambridge Univ Pr, 2007.

[30] A. Majda and X. Wang. *Non-linear dynamics and statistical theories for basic geophysical flows*. Cambridge Univ Pr, 2006.